

Product Manifold Filter: Non-Rigid Shape Correspondence via Kernel Density Estimation in the Product Space

Matthias Vestner
Technical University Munich

Roe Litman
Tel-Aviv University

Emanuele Rodolà
USI Lugano

Alex Bronstein
Technion, Israel Institute of Technology
Perceptual Computing Group, Intel, Israel

Daniel Cremers
Technical University Munich

Abstract

Many algorithms for the computation of correspondences between deformable shapes rely on some variant of nearest neighbor matching in a descriptor space. Such are, for example, various point-wise correspondence recovery algorithms used as a post-processing stage in the functional correspondence framework. Such frequently used techniques implicitly make restrictive assumptions (e.g., near-isometry) on the considered shapes and in practice suffer from lack of accuracy and result in poor surjectivity. We propose an alternative recovery technique capable of guaranteeing a bijective correspondence and producing significantly higher accuracy and smoothness. Unlike other methods our approach does not depend on the assumption that the analyzed shapes are isometric. We derive the proposed method from the statistical framework of kernel density estimation and demonstrate its performance on several challenging deformable 3D shape matching datasets.

1. Introduction

Estimating the correspondence between 3D shapes is among the fundamental problems in computer vision, geometry processing and graphics with a wide spectrum of applications ranging from 3D scene understanding to texture mapping and animation. Of particular interest is the case in which the objects are allowed to deform non-rigidly. In this setting, research has mainly focused on minimizing a measure of distortion between the input shapes, reaching in recent years very high levels of accuracy [19]. However, point-wise accuracy often comes under restricting requirements (isometry assumption), or at the price of a lack of useful properties on the computed map, namely *bijectivity* (each point on either shape should have exactly one corresponding point on the other) and *smoothness* (nearby points should match to nearby points).




Figure 1. Our method can be used to recover a dense, smooth, bijective correspondence between highly non-isometric shapes from minimal input information. In this example, we initialize our algorithm with just two hand-picked matches (red spheres on tail and front leg). Correspondence quality is visualized by transferring colors from horse to elephant via the recovered map.

In this paper, we introduce a novel method to recover smooth bijective maps between deformable shapes. Contrarily to previous approaches, we do not rely on the assumption that the two shapes are isometric. We phrase our matching problem by using the language of statistical inference, whereas the input to our algorithm is either 1) a sparse collection of point-wise matches (as few as two) which are used as landmark constraints to recover the complete map, or 2) a dense, noisy, possibly non-surjective and non-smooth map which is converted to a better map with higher accuracy and the aforementioned properties.

1.1. Related works

A traditional approach to correspondence problems is finding a *point-wise* matching between (a subset of) the points on two or more shapes. *Minimum-distortion methods* establish the matching by minimizing some structure distortion, which can include similarity of local features [33, 14, 8, 44], geodesic [30, 13, 15] or diffusion distances [17], or a combination thereof [41].

Typically, the computational complexity of such meth-

ods is high, and there have been several attempts to alleviate the computational complexity using hierarchical [37] or subsampling [40] methods. Several approaches formulate the correspondence problem as quadratic assignment and employ different relaxations thereof [42, 25, 35, 2, 15, 20]. Algorithms in this category typically produce guaranteed bijective correspondences between a sparse set of points, or a dense correspondence suffering from poor surjectivity.

Embedding methods try to exploit some assumption on the correspondence (e.g. approximate isometry) in order to parametrize the correspondence problem with a few degrees of freedom. Elad and Kimmel [18] used multi-dimensional scaling to embed the geodesic metric of the matched shapes into a low-dimensional Euclidean space, where alignment of the resulting “canonical forms” is then performed by simple rigid matching (ICP) [16, 10]. The works of [28, 38] used the eigenfunctions of the Laplace-Beltrami operator as embedding coordinates and performed matching in the eigenspace. Lipman *et al.* [26, 21, 22] used conformal embeddings into disks and spheres to parametrize correspondences between homeomorphic surfaces as Möbius transformations. Despite their overall good performance, the majority of the matching procedures performed in the embedding space often produces noisy correspondences at fine scales, and suffers from poor surjectivity. More recently, in [6, 5] the authors obtain a bijective correspondence by first computing compatible embeddings of the two shapes, and then aligning the embeddings through the use of sparse input correspondences. As opposed to point-wise correspondence methods, *soft correspondence* approaches assign a point on one shape to more than one point on the other. Several methods formulated soft correspondence as a mass-transportation problem [29, 39]. Ovsjanikov *et al.* [32] introduced the *functional correspondence* framework, modeling the correspondence as a linear operator between spaces of functions on two shapes, which has an efficient representation in the Laplacian eigenbases. This approach was extended in several follow-up works [34, 23, 3]. A point-wise map is typically recovered from a low-rank approximation of the functional correspondence by a matching procedure in the representation basis, which also suffers from poor surjectivity. A third class of matching methods formulates the correspondence problem as an optimization problem in the *product space* of the considered shapes. Windheuser *et al.* [43] seek for a two-dimensional minimal surface in the four-dimensional product space of the two input surfaces; this was later extended to a 2D-to-3D setting by Löhner *et al.* [24]. Making use of the graph structure of the considered shapes, the discretization leads to an integer linear program on the product mesh where desirable properties of the matching such as smoothness and surjectivity become linear constraints. However, the computational complexity is prohibitive even for a modestly-sized problem.




Figure 2. The Product Manifold Filter (PMF) can be applied to a variety of problems that are aiming for bijective, smooth mappings between metric spaces. Here we map a 2D shape (contour) to a 3D shape. We initialize the PMF with two semantically meaningful matches (black spheres) and obtain a dense semantically meaningful bijection.

1.2. Main contributions

Many of the works mentioned above provide a matching that is neither bijective nor smooth. In some cases the matching is only available as a sparse set of points in the product space of the two shapes. We treat these matchings as corrupted versions of the latent correspondence and propose the *Product Manifold Filter (PMF)*, a framework that increases the quality of the input mapping. We show that the considered filter leads to a linear assignment problem (LAP) guaranteeing bijective correspondence between the shapes. Despite the common wisdom, we demonstrate that the problem is efficiently solvable for relatively densely sampled shapes by means of the well-established auction algorithm [9] and a simple multi-scale approach. Unlike many of the previously mentioned techniques that assume the shapes to be (nearly) isometric, we allow them to undergo more general deformations (Figure 1) or even have different dimensionality (Figure 2).

Finally, we present a significant amount of empirical evidence that the proposed smoothing procedure consistently improves the quality of the input correspondence coming from different algorithms, including point-wise recovery methods from functional map pipelines. We also show the performance of PMF as an interpolator of sparse input correspondences.

2. A probabilistic framework

We consider a pair of three-dimensional shapes that are represented by their boundaries \mathcal{X} and \mathcal{Y} , two-dimensional manifolds embedded in \mathbb{R}^3 and thus equipped with intrinsic metrics $d_{\mathcal{X}}$ and $d_{\mathcal{Y}}$. Our goal is to find a semantically meaningful correspondence between \mathcal{X} and \mathcal{Y} . A correspondence is a *diffeomorphism* $\pi : \mathcal{X} \rightarrow \mathcal{Y}$, i.e., a smooth mapping with a smooth inverse. We do not make any other assumptions such as isometry. The correspondence π can be represented as a two-dimensional manifold Π in the four-dimensional product space $\mathcal{X} \times \mathcal{Y}$: a pair (x, y) belongs to Π iff $\pi(x) = y$. We henceforth assume that the true correspon-

dence π between \mathcal{X} and \mathcal{Y} and the manifold Π representing it are latent.

Let $\{(x_k, y_k)\}_{k \in \mathcal{K}} \subset \Pi$ be a possibly sparse sample of the said manifold. For example, these can be pairs of corresponding points on \mathcal{X} and \mathcal{Y} computed using a feature detector followed by descriptor matching. In practice, we only have access to a noisy realization of these points, $\{(\xi_k, \eta_k)\}_{k \in \mathcal{K}}$, which we assume to admit a separable i.i.d. Gaussian density, $f(\xi_k, \eta_k) \propto K(d_{\mathcal{X}}(x_k, \xi_k))K(d_{\mathcal{Y}}(y_k, \eta_k))$, where

$$K(d) = \exp\left(-\frac{d^2}{2\sigma^2}\right)$$

is an unnormalized Gaussian kernel with the parameter σ^2 . Note that the density on the manifolds is expressed in terms of the intrinsic metrics $d_{\mathcal{X}}$ and $d_{\mathcal{Y}}$.

Given the set of noisy corresponding points $\{(\xi_k, \eta_k)\}_{k \in \mathcal{K}}$ as the input, our goal is to produce a faithful estimate of the correspondence π . We propose to estimate the latent manifold Π via kernel density estimation in the product space $\mathcal{X} \times \mathcal{Y}$. To that end, we estimate the density function using the Parzen sum

$$f(x, y) \propto \sum_{k \in \mathcal{K}} K(d_{\mathcal{X}}(x, \xi_k)) K(d_{\mathcal{Y}}(y, \eta_k)). \quad (1)$$

For every point $x \in \mathcal{X}$, an estimate of $\pi(x)$ is given by a point y maximizing $f(x, y)$,

$$\hat{\pi}(x) = \arg \max_y f(x, y). \quad (2)$$

One can further impose bijectivity of $\hat{\pi} : \mathcal{X} \rightarrow \mathcal{Y}$ as a constrain, obtaining the following estimator of the entire map

$$\hat{\pi} = \arg \max_{\hat{\pi} : \mathcal{X} \rightarrow \mathcal{Y}} \int_{\mathcal{X}} f(x, \hat{\pi}(x)) dx. \quad (3)$$

The process can be iterated as shown in the one-dimensional illustration in Figure 3.

Procedures (2) or (3) have an area reduction effect on the manifold Π producing a more regular version thereof and thus a more regular correspondence π . We interpret (3) as a filter of correspondences and will henceforth refer to it a *product manifold filter* (PMF). While we defer the rigorous proof of the area reduction property to the extended version of the paper, in what follows, we illustrate it by a simple one-dimensional example.

One dimensional illustration. Let us consider a configuration of three points $\{x_-, x, x_+\}$ and the corresponding noisy points $\{y_-, y, y_+\}$ on a pair of one-dimensional manifolds \mathcal{X} and \mathcal{Y} like those depicted in Figure 3. We assume that the points are directly given in arclength parametrization, such that $d_{\mathcal{X}}(x, x_{\pm}) = |x - x_{\pm}| = b$, $d_{\mathcal{Y}}(y, y_{\pm}) =$

$|y - y_{\pm}| = a$, and $d_{\mathcal{Y}}(y, y_+) = |y - y_+| = a + \delta$. For convenience, we henceforth denote $x = y = 0$, $x_{\pm} = \pm b$, $y_- = -a$ and $y_+ = a + \delta$. In this setting, the one-dimensional manifold Π_0 representing the input correspondence in the product space comprises two segments connecting $(-b, -a)$, $(0, 0)$, and $(b, a + \delta)$, and its length is given by $L(\Pi_0) = \sqrt{b^2 + a^2} + \sqrt{b^2 + (a + \delta)^2}$.

PMF maximizes the density function

$$\begin{aligned} h(\hat{y}) &= f(0, \hat{y}) = K(0)K(\hat{y}) + \\ &\quad K(b)K(\hat{y} + a) + K(b)K(\hat{y} - a - \delta) \\ &= K(\hat{y}) + K(b)(K(\hat{y} + a) + K(\hat{y} - a - \delta)) \end{aligned} \quad (4)$$

over the values \hat{y} for the point y . First, we observe that since $K(b) > 0$, the global maximum of $h(\hat{y})$ has to be around $\hat{y} = 0$. For $\hat{y} = 0$ and $\delta = 0$, one has

$$\frac{dh}{d\hat{y}} = K'(0) + K(b)(K'(a) + K'(-a))$$

and

$$\frac{d^2h}{d\hat{y}^2} = K''(0) + K(b)(K''(a) + K''(-a))$$

Since $K'(0) = 0$ and $K'(-a) = -K'(a)$, the first derivative vanishes, while the fact that $K''(0) < 0$ and $K''(-a) = K''(a)$ implies that $\hat{y} = 0$ is the maximum of h .

Next, we perform perturbation analysis of the above maximizer by invoking the first-order Taylor expansion of h around $(\delta, \hat{y}) = (0, 0)$:

$$\frac{\partial h}{\partial \hat{y}} \approx \left. \frac{\partial h}{\partial \hat{y}} \right|_{\hat{y}=0, \delta=0} + \hat{y} \left. \frac{\partial^2 h}{\partial \hat{y}^2} \right|_{\hat{y}=0, \delta=0} + \delta \left. \frac{\partial^2 h}{\partial \hat{y} \partial \delta} \right|_{\hat{y}=0, \delta=0}.$$

Demanding equality to zero yields the maximizer of the perturbed problem

$$\hat{y} \approx \frac{K(b)K''(a)\delta}{2K(b)K''(a) + K''(0)} = \frac{\delta}{2 + \frac{K''(0)}{K(b)K''(a)}} = c\delta.$$

For $a < \frac{\sigma}{\sqrt{2}}$ the ratio in the denominator is positive and consequently $c \in (0, \frac{1}{2})$.

The length of the estimated manifold $\hat{\Pi}$ can be obtained using a series of first-order Taylor approximations,

$$\begin{aligned} L(\hat{\Pi}) &= \sqrt{b^2 + (a + c\delta)^2} + \sqrt{b^2 + (a + \delta - c\delta)^2} \\ &\approx L(\Pi_0) + \frac{ac\delta}{\sqrt{b^2 + a^2}} - \frac{(a + \delta)c\delta}{\sqrt{b^2 + (a + \delta)^2}} \\ &\approx L(\Pi_0) - \frac{cb^2}{(b^2 + a^2)^{3/2}} \delta^2 < L(\Pi_0), \end{aligned} \quad (5)$$

which manifests the length reducing effect of the PMF.

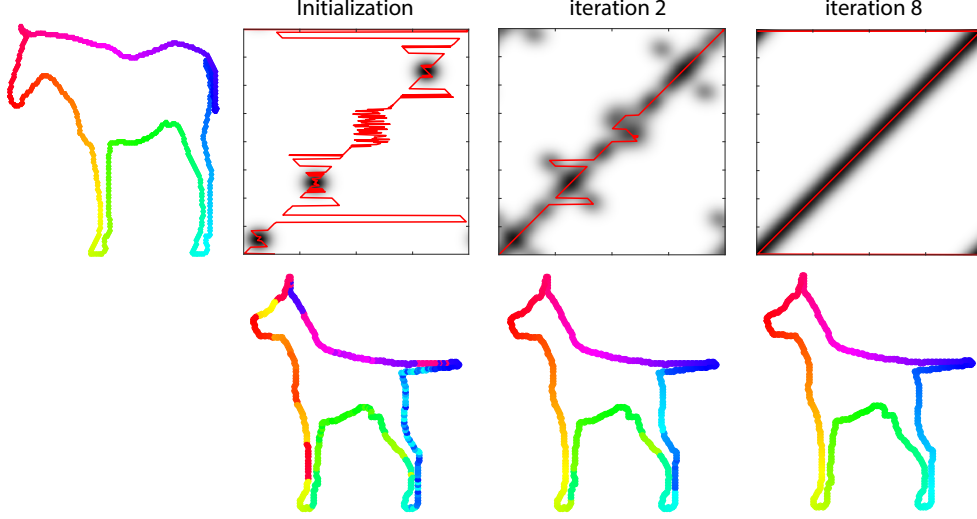


Figure 3. Conceptual illustration of our method on one-dimensional manifolds. Shown are iterations of PMF ($|\mathcal{K}| = 3$ sparse matches as initialization). Top: Kernel density estimation $f(x, y)$ as defined in (1) in the product space of the two shapes \mathcal{X} and \mathcal{Y} . Dark areas correspond to higher density. According to (3), consistently maximizing $f(x, \cdot)$ gives a bijective and smoothed matching (red curve in product space) which is used to derive the density estimate in the next iteration. Bottom: matching visualized via color transfer. Shapes are parametrized counter-clockwise with the origin of the product space corresponding to the noses of horse and dog. Note the circular boundary conditions of the product space.

2.1. Discretization

In what follows, we consider a discretization of problem (3). We assume the shape \mathcal{X} to be discretized at n points $\{x_i\}_{i=1}^n$ and the pairwise geodesic distances are stored in the matrix $\mathbf{D}_{\mathcal{X}} \in \mathbb{R}^{n \times n}$. Similarly, the shape \mathcal{Y} is discretized as $\{y_i\}_{i=1}^n$ and its pairwise distance matrix is denoted by $\mathbf{D}_{\mathcal{Y}} \in \mathbb{R}^{n \times n}$. Given a (possibly sparse) collection of input correspondences $\{(\xi_k, \eta_k)\}_{k=1}^m$ the unnormalized kernel density estimation can be written as an $n \times n$ matrix

$$\mathbf{F} = \mathbf{K}_{\mathcal{X}} \mathbf{K}_{\mathcal{Y}}^T \quad (6)$$

with the matrices $\mathbf{K}_{\mathcal{X}} \in \mathbb{R}^{n \times m}$ and $\mathbf{K}_{\mathcal{Y}} \in \mathbb{R}^{n \times m}$ given by

$$(\mathbf{K}_{\mathcal{X}})_{ik} = K(d_{\mathcal{X}}(x_i, \xi_k)) \quad (7)$$

$$(\mathbf{K}_{\mathcal{Y}})_{ik} = K(d_{\mathcal{Y}}(y_i, \eta_k)). \quad (8)$$

The objective in (3) thus becomes

$$\begin{aligned} \int_{\mathcal{X}} f(x, \pi(x)) dx &= \int_{\mathcal{X} \times \mathcal{Y}} f(x, y) \delta_{\pi(x)}(y) dy dx \\ &\approx \sum_{i,j=1}^n \mathbf{F}_{ij} \mathbf{P}_{ji} = \langle \mathbf{P}, \mathbf{F} \rangle \end{aligned} \quad (9)$$

with $\mathbf{P} \in \{0, 1\}^{n \times n}$ being a permutation matrix representing a bijection between $\{x_i\}_{i=1}^n$ and $\{y_i\}_{i=1}^n$. At some points it will be convenient to use the vector representation $p \in \{1, \dots, n\}^n$ of \mathbf{P} . Estimating the bijective correspondence as in (3) thus turns out to be a linear assignment problem (LAP) of the form

$$\hat{\mathbf{P}} = \arg \max_{\mathbf{P}} \langle \mathbf{P}, \mathbf{F} \rangle \quad (10)$$

where the optimization is performed over the space of all $n \times n$ permutation matrices.

2.2. Multiscale

While linear assignment problems like (10) can be solved in polynomial time, the memory consumption is quadratic in the vertex set size n . To alleviate this burden, we propose a multi-scale technique based on the assumption of local regularity of the manifold Π .

Given two shapes discretized at n points each, we perform farthest point sampling to obtain a hierarchy of p multiscale representations consisting of $n_1 < n_2 < \dots < n_p = n$ points. Each of the samplings comes with a sequence of sampling radii, $r_1^{\mathcal{X}} > r_2^{\mathcal{X}} > \dots > r_p^{\mathcal{X}}$ and $r_1^{\mathcal{Y}} > r_2^{\mathcal{Y}} > \dots > r_p^{\mathcal{Y}}$, respectively.

For sufficiently large shapes, the $n \times n$ pairwise distance matrices $\mathbf{D}_{\mathcal{X}}$ and $\mathbf{D}_{\mathcal{Y}}$ can be no more stored entirely in memory. We follow [4, 27] and store only the projection of the latter matrices on the first r eigenfunctions of the Laplacian resulting in an $n \times r$ matrix. The original distances are reconstructed on-demand, with negligible error as shown in [1, 27].

We recursively apply a variant of the PMF to the sparse set of input matches obtained by the coarser scale:

$$\mathbf{P}_{i+1} = \arg \max_{\mathbf{P} \in \{0,1\}^{n_{i+1}^2}} \langle \mathbf{P}, \mathbf{F}_i \rangle \quad (11)$$




Figure 4. Our method finds smooth bijective maps between non-isometric shapes even when one single match is given as input (marked as small red spheres). Note that the map remains smooth even if the initial match is wrong (rightmost column).

where

$$\mathbf{F}_i(s, t) = \sum_{k=1}^{n_i} \mathbf{W}(s, t, k, p_i(k)) \mathbf{K}_{\mathcal{X}}(s, k) \mathbf{K}_{\mathcal{Y}}(t, p_i(k)). \quad (12)$$

The weighting matrix $\mathbf{W} \in \{0, 1\}^{(n_{i+1} \times n_i)}$ assures that the image of a point x_s being in the vicinity of x_k^i is constrained to be mapped to a point in the vicinity of $p(k)$ and vice versa (i.e., the matching and its inverse are supposed to be smooth):

$$\mathbf{W}(s, t, k, l) = \begin{cases} 0 & \text{if } \mathbf{D}_{\mathcal{X}}(s, k) < r_i^{\mathcal{X}} \text{ and } \mathbf{D}_{\mathcal{Y}}(t, l) > 2r_i^{\mathcal{Y}} \\ 0 & \text{if } \mathbf{D}_{\mathcal{Y}}(t, l) < r_i^{\mathcal{Y}} \text{ and } \mathbf{D}_{\mathcal{X}}(s, k) > 2r_i^{\mathcal{X}} \\ 1 & \text{otherwise} \end{cases} \quad (14)$$

This construction leads to a sparse payoff matrix corresponding to a smaller space of feasible permutations, so that the corresponding LAP can be solved efficiently. Note the factor 2 in (14). Since we cannot guarantee the Voronoi cells on the two shapes to have the same number of points and want to be able to remove errors from the coarser scale, we permit moving a point to an adjacent Voronoi cell.

3. Experiments

While our method can be applied to a variety of problems aiming at bijective and smooth mappings between metric spaces (see Figure 2 for an extreme case), here we focus on the recovery of a correspondence between non-rigid and possibly non-isometric 3D shapes. We show the performance of our method in two very different scenarios, namely refinement of noisy dense correspondence, and completion of sparse correspondence. We additionally demonstrate the performance of our multi-scale technique by recovering bijective correspondences between high resolution shapes.

3.1. Recovery from sparse correspondences

In our first set of experiments we consider a scenario in which the input shapes come with a (possibly very sparse) collection of initial matches. These, in turn, can be obtained by a sparse non-rigid matching technique such as [35] or be hand-picked, depending on the application. In these experiments we compare PMF with the Tutte embedding approach recently introduced in [5]. Similarly to PMF, this approach produces guaranteed bijective and smooth maps starting from a sparse set of point-wise matches; to our knowledge, this method represents the state-of-the-art for this class of problems.

The results of this comparison are shown in Figure 5. The input matches were obtained by mapping farthest point samples on a reference shape via the ground-truth correspondence to the target shape, and are visualized by transferring a texture from reference to target via the recovered *dense* map. As we can read from the plots, our approach yields maps of better quality when fewer than ten matches are provided as the input, and maps of comparable quality when more matches are available. It is important to note that while our method still produces meaningful solutions when just one or two matches are given as the input (see Figures 1, 4), the approach of [5] has the theoretical minimum of three matches; furthermore, the latter approach gives different solutions depending on the specific ordering of the inputs, while our method is invariant to their permutations.

3.2. Recovery from noisy input

In this set of experiments we assume to be given a low-rank approximation of the latent correspondence \mathbf{P} in terms of a functional map

$$\mathbf{C} = \mathbf{\Psi}^T \mathbf{P} \mathbf{\Phi} \in \mathbb{R}^{r \times r}, \quad (15)$$

where $\mathbf{\Phi}, \mathbf{\Psi} \in \mathbb{R}^{n \times r}$ are truncated orthonormal bases on \mathcal{X} and \mathcal{Y} . We refer the reader to the original paper [32] for details and allow ourselves to condense its ideas to the above equation.

While a plurality of methods for finding \mathbf{C} have been proposed in the last years, there currently exist only three approaches to recover a pointwise correspondence matrix \mathbf{P} from it. In [32] the authors proposed to recover a pointwise correspondence between \mathcal{X} and \mathcal{Y} by solving the nearest-neighbor problem (NN)

$$\min_{\mathbf{P} \in \{0,1\}^{n \times n}} \|\mathbf{C} \mathbf{\Phi}^T - \mathbf{\Psi}^T \mathbf{P}\|_F^2 \quad \text{s.t. } \mathbf{P}^T \mathbf{1} = \mathbf{1}. \quad (16)$$

alternated with an orthogonality-enforcing refinement of \mathbf{C} (ICP). A variant is its bijective version (Bij. NN)

$$\min_{\mathbf{P} \in \{0,1\}^{n \times n}} \|\mathbf{C} \mathbf{\Phi}^T - \mathbf{\Psi}^T \mathbf{P}\|_F^2 \quad \text{s.t. } \mathbf{P}^T \mathbf{1} = \mathbf{1}, \mathbf{P} \mathbf{1} = \mathbf{1}. \quad (17)$$




Figure 5. Comparison between our method and the method of [5] at increasing number of input matches (reported on top). Both methods produce smooth, guaranteed bijective solutions; our method requires little computational effort (a few minutes as opposed to ~ 1 hour for [5]), and yields in comparison more accurate solutions when fed with a very sparse input.

The orthogonal refinement of (16) assumes the underlying map to be area-preserving [32], and is therefore bound to fail in case the two shapes are non-isometric. Rodolà *et al.* [36] proposed to consider the non-rigid counterpart for a given \mathbf{C} :

$$\begin{aligned} \min_{\mathbf{P} \in [0,1]^{n \times n}} \quad & D_{\text{KL}}(\mathbf{C}\Phi^T, \Psi^T\mathbf{P}) + \lambda \|\Omega(\mathbf{C}\Phi^T - \Psi^T\mathbf{P})\|^2 \\ \text{s.t.} \quad & \mathbf{P}^T \mathbf{1} = \mathbf{1}. \end{aligned} \quad (18)$$

Here D_{KL} denotes the Kullback-Leibler divergence between probability distributions, Ω is a low-pass operator promoting smooth velocity vectors, and $\lambda > 0$ controls the regularity of the assignment. The problem is then solved via expectation-maximization by the coherent point drift algorithm (CPD) [31].

We construct the low-rank functional map using the known groundtruth correspondences between the shapes. Since this is supposed to be the ideal input for all the competing methods, we abandon the refinement step in (16). Correspondences returned by the other methods are treated as noisy realizations of the latent bijection and are recovered via PMF with σ^2 set to 2% of the target shape area.

We show quantitative comparisons on 71 pairs from the SCAPE dataset [7] (near isometric, 1K vertices) and 100 pairs from the FAUST dataset [11] (including inter-class pairs, 7K vertices). In Figures 10 and 8 we compare the correspondence accuracy, while in Figure 11 we visualize how lack of smoothness, bijectivity and accuracy affect texture transfer.

The accuracy of all input-matchings is increased by applying the product manifold filter. To our knowledge the matchings obtained by the PMF are the most accurate ones that can be recovered from this type of low-rank approximation. While linear assignment problems are known to be time demanding to solve for larger numbers of variables, the most dramatic increase of run time occurs when applying the coherent point drift algorithm (see Table 1). Memory



Figure 6. Result of our method on two cat shapes from TOSCA [12]. This high resolution shape has 27894 vertices, making it infeasible to store the entire pairwise distance matrix in memory. Our multi-scale approach recovered a smooth matching from only 20 sparse correspondences given as the input using five hierarchical scales as detailed in Section 2.2.

consumption and run times however limit the direct applicability of the PDM.

3.3. Recovering high-resolution correspondences using multiscale

In this set of experiments we demonstrate how the PMF together with the multiscale method described in 2.2 can recover very accurate matchings on shapes being sampled at high resolution. Figure 6 shows a dense bijective matching between two shapes sampled at $n = 27894$ points each. At each of the six scales $n_i \in \{10^3, 2 \times 10^3, 4 \times 10^3, 8 \times 10^3, 1.6 \times 10^4, n\}$ the constrained LAP (11) was solved. Figure 7 shows the improvement of correspondence accuracy at finer scales. By using the weighting functions we force points to stay close to their nearest neighbor in the coarser sampling and thus can guarantee to approximately keep the accuracy from the coarser scale. Solving the constrained LAP at the finest resolution took less than 9 minutes. Calculating the kernel density matrices (6) for all




Figure 7. Quantitative analysis of correspondences between the two cats shown in Figure 6, recovered using the multiscale approach. The geodesic errors are measured with respect to the groundtruth on the finest scale. At coarse scales the minimal expected geodesic error introduced by any matching is in the order of the sampling radius. As expected, the accuracy of the matching increases with each iteration.

n	1000	1000	6890	6890
r	20	50	20	50
Nearest neighbors	0.04	0.06	1.35	2.88
Bijjective NN	2.79	2.30	463.66	253.03
CPD	4.79	4.67	1745.06	2085.65
NN + PMF	1.75	1.28	382.86	244.10
Bij. NN + PMF	4.06	3.44	746.00	440.94

Table 1. Average runtimes in seconds. We compare the runtimes of different correspondence recovery methods. Given the rank k of a functional map approximating the correspondence between shapes sampled at n points each, we report the time it takes to obtain a dense matching. See Figures 10,8 and 11 for evaluations of accuracy.

scales took less than 40 minutes.

Another test was performed on pairs of shapes from the FAUST dataset. As Figure 9 shows, the correspondences obtained using a four-scale scheme are comparable in accuracy to the solution of a single-scale scheme. However, the runtime of the multi-scale approach is significantly lower. Calculating the kernel density matrices (6) for all scales took about 4 minutes, while solving the LAPs at all scales took around 18 seconds.

4. Discussion and conclusion

We considered the problem of bijective correspondence recovery by means of filtering a given set of matches coming from any of the existing algorithms (including those not guaranteeing bijectivity, or producing sparse correspondences). Viewing correspondence computation as a kernel density estimation problem in the product space, we introduced the product manifold filter that leads to smooth cor-




Figure 8. Quantitative comparison of methods for pointwise correspondence on the non-isometric FAUST dataset (about 7K vertices).




Figure 9. Error curves for a multiscale experiment on the FAUST dataset, showing result for intermediate scales. For comparison, the solution obtained by a single-scale PMF is shown in dashed black.

respondences, with the additional constraint of bijectivity embodied through an LAP. We believe that statistical tools that have been heavily used in other domains of science and engineering might be very useful in shape analysis, and invite the community to further explore this direction. Of special interest is the possibility to lift the product space to higher dimensions encoding local similarity of points on the two shapes, for instance by using descriptors. The way the kernel density estimator is constructed does not restrict the samples per shape to be distinct. Together with the use of weighting factors this allows to directly work with softmaps as inputs. Finally, we believe that denoising the correspondence manifold in the product space is a useful perspective applicable to different problems in computer vision where smooth correspondences are desired, such as optical flow.

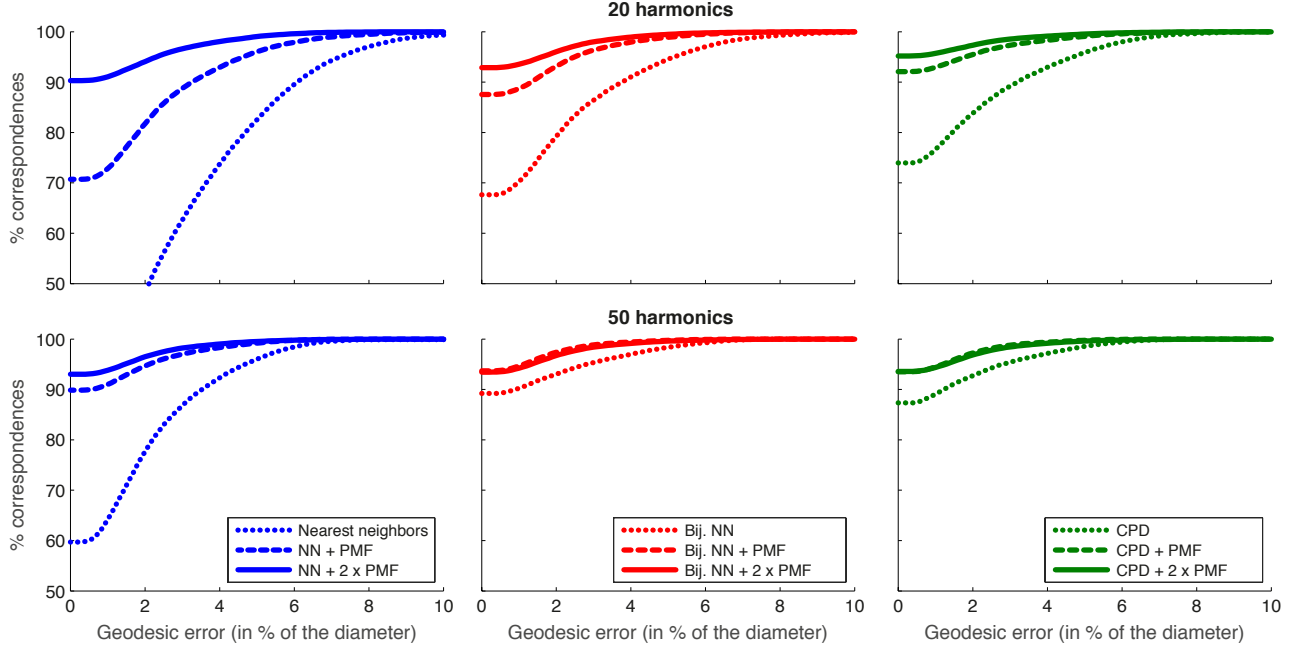


Figure 10. Quantitative comparison of methods for pointwise correspondence recovery from a functional map (20 and 50 eigenfunctions). We matched 70 pairs from the near-isometric SCAPE dataset (1K). Plotted are the histograms of geodesic errors. Filtering the results of nearest neighbors (left) outperforms the state of the art method (right) while having only a fraction of its runtime (Table 1). Even better results are achieved under affordable runtimes when initializing PMF estimator with the result of bijective NN (center).




Figure 11. Qualitative comparison of methods for pointwise correspondence recovery from a functional map. Current methods such as Nearest Neighbors (NN) and coherent point drift (CPD) suffer from bad accuracy and lack of surjectivity. Applying the proposed estimation to either of them gives a guaranteed bijective matching with high accuracy and improved smoothness. Left: We visualize the accuracy of the methods by transferring texture from the source shape \mathcal{X} to the target shape \mathcal{Y} . Neither Nearest Neighbors nor CPD produce bijective mappings. The lack of surjectivity is visualized by assigning a fixed color (green) to not-hit points. Right: The geodesic error (distance between groundtruth and recovered match, relative to the shape diameter) induced by the matching is visualized on the target shape \mathcal{Y} .

References

- [1] Y. Aflalo, H. Brezis, and R. Kimmel. On the optimality of shape and data representation in the spectral domain. *SIAM Journal on Imaging Sciences*, 8(2):1141–1160, 2015. 4
- [2] Y. Aflalo, A. Bronstein, and R. Kimmel. On convex relaxation of graph isomorphism. *PNAS*, 112(10):2942–2947, 2015. 2
- [3] Y. Aflalo, A. Dubrovina, and R. Kimmel. Spectral generalized multidimensional scaling. *IJCV*, 2016. 2
- [4] Y. Aflalo and R. Kimmel. Spectral multidimensional scaling. *PNAS*, 110(45):18052–18057, 2013. 4
- [5] N. Aigerman and Y. Lipman. Hyperbolic orbifold tutte embeddings. *ACM Transactions on Graphics (TOG)*, 35(6):217, 2016. 2, 5, 6
- [6] N. Aigerman, R. Poranne, and Y. Lipman. Lifted bijections for low distortion surface mappings. *ACM Transactions on*

- Graphics (TOG)*, 33(4):69, 2014. 2
- [7] D. Anguelov et al. SCAPE: Shape completion and animation of people. *TOG*, 24(3):408–416, 2005. 6
- [8] M. Aubry, U. Schlickewei, and D. Cremers. The wave kernel signature: A quantum mechanical approach to shape analysis. In *Proc. ICCV*, 2011. 1
- [9] D. P. Bertsekas. *Network optimization: continuous and discrete models*. Citeseer, 1998. 2
- [10] P. J. Besl and N. D. McKay. A method for registration of 3D shapes. *PAMI*, 14(2):239–256, 1992. 2
- [11] F. Bogo, J. Romero, M. Loper, and M. J. Black. FAUST: Dataset and evaluation for 3D mesh registration. In *Proc. CVPR*, 2014. 6
- [12] A. Bronstein, M. Bronstein, and R. Kimmel. *Numerical Geometry of Non-Rigid Shapes*. Springer, 2008. 6
- [13] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *PNAS*, 103(5):1168–1172, 2006. 1
- [14] M. M. Bronstein and I. Kokkinos. Scale-invariant heat kernel signatures for non-rigid shape recognition. In *Proc. CVPR*, 2010. 1
- [15] Q. Chen and V. Koltun. Robust nonrigid registration by convex optimization. In *Proc. ICCV*, 2015. 1, 2
- [16] Y. Chen and G. Medioni. Object modeling by registration of multiple range images. In *Proc. Conf. Robotics and Automation*, 1991. 2
- [17] R. R. Coifman, S. Lafon, A. B. Lee, M. Maggioni, B. Nadler, F. Warner, and S. W. Zucker. Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps. *PNAS*, 102(21):7426–7431, 2005. 1
- [18] A. Elad and R. Kimmel. Bending invariant representations for surfaces. In *Proc. CVPR*, 2001. 2
- [19] O. Kaick, H. Zhang, G. Hamarneh, and D. Cohen-Or. A survey on shape correspondence. *Computer Graphics Forum*, 20:1–23, 2010. 1
- [20] I. Kezurer, S. Kovalsky, R. Basri, and Y. Lipman. Tight relaxations of quadratic matching. *Computer Graphics Forum*, 34(5), 2015. 2
- [21] V. G. Kim, Y. Lipman, X. Chen, and T. A. Funkhouser. Möbius transformations for global intrinsic symmetry analysis. *Computer Graphics Forum*, 29(5):1689–1700, 2010. 2
- [22] V. G. Kim, Y. Lipman, and T. Funkhouser. Blended intrinsic maps. *TOG*, 30(4):79, 2011. 2
- [23] A. Kovnatsky, M. M. Bronstein, X. Bresson, and P. Vandergheynst. Functional correspondence by matrix completion. In *Proc. CVPR*, 2015. 2
- [24] Z. Löhner, E. Rodolà, F. R. Schmidt, M. M. Bronstein, and D. Cremers. Efficient globally optimal 2d-to-3d deformable shape matching. In *Proc. CVPR*, 2016. 2
- [25] M. Leordeanu and M. Hebert. A spectral technique for correspondence problems using pairwise constraints. In *Proc. ICCV*, 2005. 2
- [26] Y. Lipman and I. Daubechies. Conformal Wasserstein distances: Comparing surfaces in polynomial time. *Advances in Mathematics*, 227(3):1047 – 1077, 2011. 2
- [27] R. Litman and A. Bronstein. Spectrometer: Amortized sub-linear spectral approximation of distance on graphs. In *2016 International Conference on 3D Vision*, 2016. 4
- [28] D. Mateus, R. P. Horaud, D. Knossow, F. Cuzzolin, and E. Boyer. Articulated shape matching using laplacian eigenfunctions and unsupervised point registration. In *Proc. CVPR*, 2008. 2
- [29] F. Mémoli. Gromov-Wasserstein Distances and the Metric Approach to Object Matching. *Foundations of Computational Mathematics*, pages 1–71, 2011. 2
- [30] F. Mémoli and G. Sapiro. A theoretical and computational framework for isometry invariant recognition of point cloud data. *Foundations of Computational Mathematics*, 5(3):313–347, 2005. 1
- [31] A. Myronenko and X. Song. Point set registration: Coherent point drift. *TPAMI*, 32(12):2262–2275, 2010. 6
- [32] M. Ovsjanikov, M. Ben-Chen, J. Solomon, A. Butscher, and L. Guibas. Functional maps: a flexible representation of maps between shapes. *ACM Trans. on Graphics*, 31(4):30, 2012. 2, 5, 6
- [33] M. Ovsjanikov, Q. Mérigot, F. Mémoli, and L. Guibas. One point isometric matching with the heat kernel. *Computer Graphics Forum*, 29(5):1555–1564, 2010. 1
- [34] J. Pokrass, A. M. Bronstein, M. M. Bronstein, P. Sprechmann, and G. Sapiro. Sparse modeling of intrinsic correspondences. In *Computer Graphics Forum*, volume 32, pages 459–468. Wiley Online Library, 2013. 2
- [35] E. Rodolà, A. M. Bronstein, A. Albarelli, F. Bergamasco, and A. Torsello. A game-theoretic approach to deformable shape matching. In *Proc. CVPR*, 2012. 2, 5
- [36] E. Rodolà, M. Moeller, and D. Cremers. Point-wise map recovery and refinement from functional correspondence. In *Proceedings Vision, Modeling and Visualization (VMV)*, Aachen, Germany, 2015. 6
- [37] Y. Sahillioğlu and Y. Yemez. Coarse-to-fine combinatorial matching for dense isometric shape correspondence. *Computer Graphics Forum*, 30(5):1461–1470, 2011. 2
- [38] A. Shtern and R. Kimmel. Matching lbo eigenspace of non-rigid shapes via high order statistics. *arXiv:1310.4459*, 2013. 2
- [39] J. Solomon, A. Nguyen, A. Butscher, M. Ben-Chen, and L. Guibas. Soft maps between surfaces. In *Computer Graphics Forum*, volume 31, pages 1617–1626, 2012. 2
- [40] A. Tevs et al. Intrinsic shape matching by planned landmark sampling. *Computer Graphics Forum*, 30(2):543–552, 2011. 2
- [41] L. Torresani, V. Kolmogorov, and C. Rother. Feature correspondence via graph matching: Models and global optimization. In *Proc. ECCV*, 2008. 1
- [42] S. Umeyama. An eigendecomposition approach to weighted graph matching problems. *PAMI*, 10(5):695–703, 1988. 2
- [43] T. Windheuser, U. Schlickewei, F. R. Schmidt, and D. Cremers. Geometrically consistent elastic matching of 3d shapes: A linear programming solution. In *IEEE International Conference on Computer Vision (ICCV)*, 2011. 2
- [44] A. Zaharescu, E. Boyer, K. Varanasi, and R. Horaud. Surface feature detection and description with applications to mesh matching. In *Proc. CVPR*, 2009. 1